DOCUMENT RESUME

ED 377 709 FL 022 680

AUTHOR Moisl, Hermann
TITLE Recurrent Artificial Neural Networks and Finite State Natural Language Processing.
PUB DATE [94]
NOTE 13p.
PUB TYPE Reports - Evaluative/Feasibility (142) -- Viewpoints (Opinion/Position Papers, Essays, etc.) (120)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Computational Linguistics; Foreign Countries; *Language Processing; Language Research; *Linguistic Theory; Neurolinguistics; Simulation
IDENTIFIERS *Natural Language; *Neural Networks

ABSTRACT

It is argued that pessimistic assessments of the adequacy of artificial neural networks (ANNs) for natural language processing (NLP) on the grounds that they have a finite state architecture are unjustified, and that their adequacy in this regard is an empirical issue. First, arguments that counter standard objections to finite state NLP on the grounds that these objections confuse the explanatory aims of linguistic theory with the essential technological aims of NLP are presented. A finite state NLP model that maps strings onto meaning representations is then proposed and preliminary string processing test and cluster analysis results of a computer simulation are presented, briefly addressing some problems and further developments. It is concluded that it remains to be seen whether the model can be developed for general NLP, but that it exemplifies the kind of radical departure from linguistics-based NLP that is possible once the supposed theoretical obstacles to finite state NLP are removed. In particular, it departs from linguistics-based NLP in making no use of any syntactic or compositional structure beyond the purely sequential and amounts to table lookup mapping from strings to meaning representations. Contains 33 references. (MSE)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 377 709

Recurrent Artificial Neural Networks and Finite State Natural Language Processing

Hermann MOISL
Department of English Language
University of Newcastle upon Tyne
Newcastle upon Tyne NE1 7RU
Great Britain

Email: Hermann.Moisl@newcastle.ac.uk

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Herman Moisl
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

Abstract

This paper argues that the supposed inadequacy of recurrent ANNs for NLP on the grounds that they have finite-state architecture is mistaken in that it confuses the explanatory aim of generative linguistic theory with the essentially technological one of NLP, which is to construct physical string-processing devices. The first part of the discussion counters standard objections to finite state NLP in the light of this assumption, and the second proposes, implements, and tests a finite state NLP model which maps strings to meaning representations.

Keywords: Natural language processing, artificial neural networks, finite-state NLP.

Introduction

One of the many application areas of artificial neural networks (ANN) has been natural language processing (NLP), and various ANN approaches to language processing have been proposed (Sharkey & Reilly 1992 for overview and recent work). This paper concentrates on one which has been prominent in recent research: induction of language processing devices in synchronous recurrent distributed ANNs by exposure to symbol string environments (Elman 1990, 1991; Giles *et al* 1990, 1991, 1992, 1993; Servan-Schreiber *et al* 1989, 1991; Pollack 1991; Sharkey & Sharkey 1993; Watrous & Kuhn 1992; Castano *et al* 1993; overview in Sanfeliu & Alquezar 1992). Such nets have yielded promising results in language processing and in NLP more particularly, but there appears to be a problem in principle which renders them inadequate for general NLP. In terms of their architecture and string processing dynamics these nets look very like strict finite state automata, where 'strict' refers to FSAs in which there is no distinction between processor and memory (on which see further Schwarz 1992). But one of the earliest and most enduring results in generative linguistics is that finite state devices are inadequate for generating natural languages, with the consequence that the nets in question appear to be inadequate on theoretical grounds for NLP. The general argument of this paper is that the theoretical problem is only apparent, and that the adequacy of recurrent distributed nets for NLP is a purely empirical issue. The discussion is in two main parts. The first considers and rejects standard arguments against finite state NLP,

FL 022 680

BEST COPY AVAILABLE

and the second sketches a general approach to finite state NLP using one particular kind of recurrent ANN: the simple recurrent network (SRN).

1. Finite state processing in NLP

NLP research has been and continues to be naturally associated with other disciplines that concern themselves with the study of human language --theoretical linguistics most obviously, but also the range of disciplines that come under the umbrella of cognitive science, for which theoretical characterization of natural language is widely seen as paradigmatic for understanding of cognition more generally. The argument of this section is based on the assumption that the aims of NLP and of these associated disciplines are fundamentally different. Mainstream linguistics and cognitive science regard natural language as an abstract object, and aim to explain it by developing maximally expressive and economical theories about it. NLP, on the other hand, is concerned with the design and construction of physical devices to process physical strings for some purpose. On the basis of this assumption, the current section argues that the objections in principle to finite state NLP stem from a failure to keep these aims separate.

The pretty much standard approach in non-ANN NLP, exemplified in Gazdar & Mellish 1989, is to implement a physical device on the basis of some preferred linguistic theory. That is, the grammatical categories, syntactic structures, and any structure-manipulating devices such as movement rules of the linguistic theory are incorporated into a language-processing algorithm, and that algorithm can then either be directly physically realized or simulated on a general-purpose computer. The result is a physical NLP device whose behaviour is both completely specified by the linguistic theory and as fully understood as the theory itself. Moreover, as a purely practical matter, when simulation rather than direct realization is chosen --and it invariably is-- the construction of the NLP device becomes no more difficult than design and coding of the algorithm; a compiler takes care of the rest. These are considerable advantages, but the fact remains that direct instantiation of linguistic theories is not a necessary approach to construction of NLP devices.

Given a physical system with an observable behaviour, what internal mechanisms produce that behaviour? This is the identification problem, and the answer, in Arbib's words, is this: 'Even if we know completely the function, or behaviour, of a device, we cannot deduce from this a unique structural description...The process of going from the behaviour of a system to its structural description is then not to be thought of as actually identifying the particular state variable form of *the* system under study, but rather that of identifying a state variable description of *a* system that will yield the observed behaviour, even though the mechanism for generating that behaviour may be different from that of the observed system' (1987, pp.38-9). True, theoretical linguistics is concerned with NL in the abstract, not with the language behaviour of individual systems/humans, but the identification problem nevertheless applies in that a class of systems is at issue: linguistic theory is a structural description of the class of natural language speakers. This applies straightforwardly to the present discussion. Assume some function, say, a mapping from sentences to meanings. An NLP device realizes the function by pairing physical sentence representations with physical meaning representations. Such a device would be describable in terms of a processor which uses linguistic-theoretic ontology even though

it did not physically instantiate that ontology but rather used some other mechanism. In other words, linguistic theory has no necessary implications for the design of NLP processors. In principle, then, the way is open for candidate NLP technologies, including devices with finite-state architecture such as recurrent ANNs. There are, however, some standard objections to finite state architectures for NLP, and these will now be dealt with in turn.

a) Unbounded-length centre embedding strings

One of the aims of syntactic theory is to characterize the range of NL sentence structures as economically as possible without any necessary regard to how those characterizations might relate to production and understanding of utterances by speakers in the real world. Some of these structures are characterized as recursive; because it imposes no physical realization constraints, syntactic theory permits an arbitrary depth of recursive embedding and consequently sentences of unbounded length, which in turn renders the language in question an infinite string set. Now, as long ago as 1959 Chomsky demonstrated that finite state devices are incapable of generating or processing the language $a^n b^n$, that is, a set of strings in which some unbounded number of a given symbol a is followed by exactly the same number of some other symbol b, where n is any positive integer. Since this string pattern is attested in NLs in what are analyzed as recursive centre embedding structures, and since generative grammars do not specify limits on recursive structures, any NL processing device has to be able to deal with unbounded centre embedded strings. But this is impossible if the device is finite state, and so finite state devices are inadequate for NL processing.

If, however, one is concerned not with abstract characterization of natural language but with constructing a physical NLP device, it is only necessary to consider the string set which the device can be expected to encounter in practice rather than the set it might encounter in principle. On that view, arguments against the finite stateness of NLs simply do not apply. In the real world there is no such thing as an arbitrarily deeply nested recursive structure, and no such thing as a string of unbounded length. These things have finite limits, and, seen in terms of what infinity is conventionally taken to mean, even the longest NL strings are really very short. The difference between the theoretical linguist's and the NLP researcher's views of NL comes down to this. A grammar with recursion generates an infinite string set as long as there is no bound on the application of that recursion. As soon as a bound is imposed the language which the grammar generates becomes a finite subset of the one generated by the unbounded version, which can be processed by a finite state device (Hopcroft and Ullman 1979). There is consequently no reason why a finite state device should not process any member of the class of natural languages.

None of this is new. Miller & Chomsky (1965, 464-83) recognized that the human language processor had to be finite state, and the inevitable finite stateness of physical NLP devices has since been generally accepted though rarely mentioned and sometimes, it seems, forgotten. Despite this, strict finite state architecture has never to my knowledge been seriously entertained for NLP. The reasons for this have to do with the second and third objections, to which we now turn.

b) Capturing generalizations

Strict finite state architecture differs from that of higher-order automata in automata theory in that the the higher-order ones have a processor-memory distinction (Schwarz 1992). If a bound is placed on the memory of such a higher-order device, it becomes functionally finite state, that is, it can be simulated in terms of input-output behaviour by a strict finite state machine. The only difference is that the two sorts of device use different algorithms to compute any given function on account of their different architectures.

Advocates of finite state NLP have preferred higher-order, bounded-memory architectures to the strict finite state one explicitly on account of the former's explanatory advantage in relation to generative linguistic theory (Chomsky 1956, 109; Miller & Chomsky 1965, 466-75; Church 1980; Pulman 1986). If one's aim is explanation, then clearly a higher-order automaton with bounded memory is to be preferred. But the primary aim of NLP is construction of physical devices, not explanation, so the choice between higher-order, bounded memory and strict finite state architecture is in principle neutral.

c) Compositionality

Semantic theory in theoretical linguistics aims to associate linguistic expressions with meanings, given some definition of 'meaning'. This involves at least:

- assignation of meaning to the primitive expressions --the morphemes-- of a given language,
- specification of how the primitive expressions relate to composite expressions -- phrases and sentences-- in that language, and
- specification of how linguistic meaning relates to the world.

There are various approaches to these tasks; the one used for exemplification here, formal or Montague semantics (Cann 1993), has been and continues to be influential. Formal semantics defines the meaning of a sentence as its truth conditions: a sentence means what the world would have to be like for the sentence to be true. Truth conditions are defined relative to a 'universe of discourse' which the semantic theory models. Such a model has two main parts: (1) the entities of the universe of discourse are identified, and the relationships between these entities and the primitive expressions of the language, that is, the denotations of the primitive expressions, are defined, and (2) how the denotations of composite expressions are constructed from those of primitive expressions is specified. Part (2) is based on Frege's principle of compositionality, which says that the meaning of a composite expression is a functions of the meanings of its component primitives and their manner of combination, and is defined using a grammar which both generates the sentences of the language in question, and associates a constituent structure with each sentence. The connection between syntactic structure and meaning is made by the rule-to-rule hypothesis, in which each syntactic rule is associated with a corresponding semantic rule which specifies the meaning of a composite expression in terms of its immediate syntactic constituents. In this way, syntax can be said to drive semantics in the sense that, given the meanings of the morphemes of a language, syntactic structure determines what sentences mean.

Compositionality has also been at the centre of a long-running debate in the cognitive science community, where the rival claims of an established 'classical' approach to the study of cognition and those of the 'connectionist' challenger are at stake (recent

discussion and extensive refererences in Dinsmore 1992 and *Connection Science* vol 4). The classical position, forcefully put by Fodor and Pylyshyn in 1988, is that, in order adequately to explain certain fundamental aspects of cognition one requires the notion of the structured representation, that is, of a mental object consisting of primitive symbols arranged in a constituent structure with compositional semantics, and the notion of mental processes which interpret representations in a way that is sensitive to their structure; because they are by nature finite state devices, ANNs cannot articulate or process representations with a constituent structure adequate for capturing the requisite generalizations about cognition. In response, connectionist cognitive scientists have worked to vindicate ANNs as a suitable alternative paradigm for cognitive theorizing by developing specifically connectionist accounts of compositionality, chief among them Smolensky's tensor product representation (Smolensky 1990) and van Gelder's functional compositionality (van Gelder 1990), which is based on temporal rather than spatial structuring of constituents. Functional compositionality is exemplified in Pollack's RAAM (Pollack 1990) architecture, and applied to NLP by Chalmers 1990 and Blank, Meeden, & Marshall 1991.

Central to both theoretical linguistics and cognitive science, then, is the idea of semantic interpretation of an expression on the basis of its syntactic structure. Now, it is universally agreed that NL sentences have a complex and varied structures. How, therefore, can a finite state architecture which imposes a single, strictly sequential syntax on sentences be adequate for NL semantics? The classicists in the cognitive science debate think that it cannot, and the connectionsts agree in that they have found it necessary to develop compositional representational methods in ANNs. Nevertheless, I claim that compositionality has no necessary implications for NLP.

The basis for this claim is, again, that explanation is one thing, and construction of physical devices another. Compositional semantics in NL sentences depends crucially on the attribution of more or less complex syntactic structure to sentences. But sentences are abstract objects, and so are the structures attributed to them: both are artefacts of linguistic theory. Physical strings, whether spoken or written, have no physical structure apart from strict temporal or spatial sequence. One approach to instantiating a sentence-to-meaning function is to map strings onto abstract sentences with abstract structures and to process these sentences in accordance with those structures, which yields a processor that is perspicuously related to linguistic theory as well as structured meaning representations, but the identification problem says that this is not necessary. Since, for NLP, it suffices that the sentence-to-meaning function be realized by a device which pairs strings with physical meaning representations, another legitimate approach is to attempt to process strings in accordance with the structure we know for certain they have --strict sequence-- using a process with strict finite state architecture.

2. Recurrent ANNs and NLP

Assuming the validity of the arguments in Part 1, the way is clear for development of an approach to NLP based on strict finite state architecture. But, as the Celtic chieftain Calgacus is reputed to have said of the Roman invasion of Britain: 'They made a desert and called it peace' (Tacitus, *Agricola*). Notions of complex phrase structure and associated compositional semantics provide intuitively accessible and theoretically well developed ways of thinking about language, and offer a basis for the design of NLP

devices via the relationship between linguistics and computation which formal language and automata theory defines. If one dispenses with these notions in NLP design, what will replace them? This section sketches a proposal for an alternative.

It is best to be clear at the outset that generative linguistic theory is not being challenged: for present purposes it is accepted as a characterization of what has traditionally been called human linguistic competence. The issue is the relationship between linguistic theory and the physical processing mechanisms required for NLP, and the argument is, in essence, that physical instantiation of computational architectures with the processor-memory distinction (Schwarz 1992) is unnecessary.

Generative linguistic theory defines a function from linguistic expressions --words, phrases, sentences-- to meanings, given some definition of 'meaning'. An NLP device implements this function if, given a physical representation of an expression, it returns a physical representation of the associated meaning. Since NLs are finite sets for NLP purposes, it becomes possible to define the function as a list of (expression-meaning) pairs, and for an NLP device to do nothing more than table lookup on a physical representation of the list. The proposal is to train a recurrent ANN with strict finite state architecture to implement such a table lookup device. The discussion is in three parts: (a) designs a computer simulation of an ANN to associate strings with meaning representations, (b) presents test results and analysis of that ANN, and (c) briefly addresses some issues which arise from (a) and (b).

a. Network design and implementation

The aim is to implement a mapping from linguistic expressions to meanings, given a finite set of expressions. This requires pairing of each expression of length 1 with a meaning, each expression of length 2 with a meaning, and so on up to some maximum length. For example:

	Meanings					
	1	2	3	4	5	6
Expressions	the					
	the	man				
	the	man	in			
	the	man	in	the		
	the	man	in	the	tent	
	the	man	in	the	tent	slept

Figure 1

Here, meanings are labelled 1, 2, 3...; the expression *the* means 1, *the man* means 2, *the man in* means 3, and so on. In view of the number of words in English and of the possible combinations of words up to some reasonable maximum sentence length, this approach may appear to require a very large number of meanings, but in fact there are exactly as many as that which generative linguistic theory would posit for the same expression set. Linguistic theory would generate the meanings more elegantly than an explicit listing by building them out of the meaning primitives assigned to morphemes, but elegance is not an issue here.

Like much of the grammatical induction work mentioned in the Introduction, the net used to implement this mapping is a simple recurrent network (SRN), a discrete-time

dynamical system whose architecture and processing dynamics make it straightforwardly interpretable as a finite state automaton (for example Arbib 1987, 24-6).

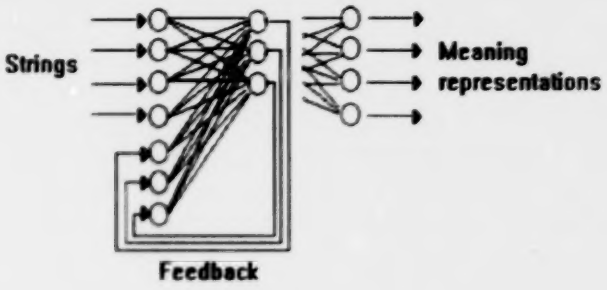

Figure 2

- The set of inputs is the input alphabet
- The set of outputs is the output alphabet
- The set of hidden layer configurations is the state set
- The connections between the input and hidden layers are the next state function in that, for every combination of current input and current state, they generate a characteristic associated next state in the hidden layer.
- The connections between hidden and output layers are the output function in that, for every state of the hidden layer, they generate a characteristic output in the output layer.

The SRN is in principle a physical device with physical input and output signals, but as with virtually all ANN research practice it is simulated on a conventional computer. The linguistic expressions and meanings, which are in principle to be represented physically, in fact need to be given a representation appropriate to the simulation. This means numerical vectors whose components represent features of the physical input and output signals. The frequently-used 'one-hot' encoding is adopted here, where each distinct input and output from the net is assigned a unique binary-valued vector in which a single component is 1 and all the rest are 0. This is unrealistic for actual NLP work, where one would want to represent features of acoustic or visual input in some detail, but it suffices for present purposes. The result is a list of (binary-valued vector sequence, binary valued vector) pairs which represent the (linguistic expression, meaning) pairs of the mapping.

The SRN is trained by repeated random selection of a pair from the list, and presentation of that pair to the net so that it can learn to associate the components using back propagation; this continues until the learning error curve stops decreasing significantly. For example, training the net on the list in Figure 1 would proceed as follows. Choose a pair, say (*the*,1): the vector representation of *the* is presented to the designated units in the first layer of the net, the vector representation of 1 becomes the target output, and back propagation is applied. Choose another, say (*the man*, 2): the vector representation of *the* becomes the input and the vector representation of 1 the target output, then the vector representation of *man* becomes the input and the vector representation of 2 the target output, applying back propagation in each case. Similarly, for (*the man in*, 3), *the* is associated with 1, *man* is associated with 2, and *in* with 3. In

this way, the net learns a mapping from strings to meanings: in all the Figure 1 strings *the* means 1, in strings 2-6 *the man* means 2, in strings 3-6 *the man in* means 3, and so on up to *the man in the tent slept*, which means 6.

b. Results and analysis

i. Results:

The simulation was trained and tested on sets of up to 24 strings of maximum length 12, all of them of the declarative *the cat with the long tail sat on the mat* variety, and learned the string-meaning function perfectly. Special attention was paid to the net's ability to handle long-distance dependencies, since this is always an issue in NLP research and has often --though erroneously-- been regarded as a problem for finite state processors. The net had no difficulty with maintaining dependencies across the distances so far required of it: *the man in the boat by the shore sees / the men in the boat by the shore see*, or the centre-embedding *the car the man the woman loves drove stopped.*

ii. Analysis:

For an SRN with n units in its hidden layer, the values which those units assume at any processing step constitute an n-component vector. Each such vector defines a point in n-space. For a test string of length a there are a hidden-layer configurations, and for b test strings there are $c = \Sigma_{1..b}\, \Sigma_{1..a}$ configurations. In the grammatical inference research mentioned in the Introduction, it is usual to carry out a cluster analysis of these c vectors in order to gain some insight into network operation. These analyses have shown that the points are not in general haphazardly distributed in n-space, but cluster in regions of that space such that all the hidden-layer vectors which generate a given output are adjacent. Such clusters are usually interpreted as states of a finite state machine which the net has inferred from string input, but this is an unnecessary abstraction. The component vectors of a cluster typically differ from one another to greater or lesser degrees, and an alternative interpretation, adopted here, is to assign a separate state to each distinct hidden layer configuration. This results in a much larger finite state machine for a given function, but corresponds more directly to physical reality.

The hidden layer configurations which the trained net assumed in the course of string processing in (i) were subjected to cluster analysis, and the results were unsurprising: the hidden-layer vectors clustered in accordance with the target output, as above. Assuming, however, a separate state for every distinct vector, the analysis shows that there is a unique state for every (expression-meaning) pair. Thus, for the three strings

1. *the man ate his lunch*
2. *the man in the restaurant ate his lunch*
3. *the man in the restaurant ate a sparing lunch*

the state sequences for initial *the* and for *man* are identical in all three cases; thereafter (1) and (2,3) bifurcate, and (2,3) continue with identical state sequences for *in the restaurant ate*, at which point they too bifurcate; *ate his lunch* in (1) and (2), though lexically identical, have different state sequences; there are three distinct states for *lunch*. By following the trajectory which each string generates in the state space, one can see that the net learns a distinct state sequence for every distinct string.

c. Discussion

Whether or not the approach proposed in (a) will be practicable, and what sorts of development are required, are empirical matters. At least two issues arise at this stage, however, both of them having to do with the assumed prior existence of an (expression, meaning) list and of a corresponding physical representation of it:

i. Though theoretically possible on the grounds of finiteness, construction of an (expression, meaning) list for some NL would be a huge and almost certainly infeasible task.

ii. Linguistic expressions and meanings are abstractions, and compilation of a list implies some physical representation scheme. Representation of linguistic expressions is well understood, but it is not at all clear what a physical representation of a meaning might look like; in the foregoing discussion 'meaning representations' were simply variables instantiated by arbitrary vectors for the sake of argument.

Despite its theoretical validity, the architecture proposed in (a) emerges as impractical in the light of (i) and (ii). The following development of that architecture is intended to address this.

An increasingly favoured idea in connectionist research is that ANN-based systems be 'hooked up' to the world with a range of sensors which transduce physical signals into a format amenable to ANN processing, thus providing input representations which relate to environmental regularities in a systematic way and thereby, it is hoped, allowing nets to behave in ways interpretable as semantically coherent (e.g. Plunkett *et al* 1992; Pfeifer & Verschure 1992; Peschl 1992). My development of the NLP approach being proposed here is based on this idea. Specifically, I use several SRNs to implement a system of interconnected finite state automata:

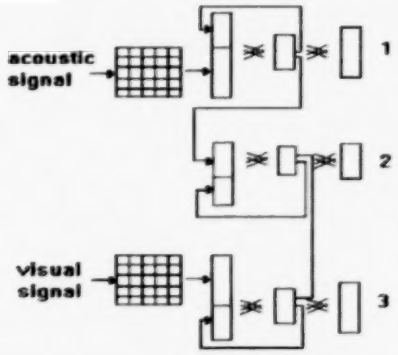

Figure 3

The acoustic and visual components (1, 3) are the input sensors. Each consists of a Kohonen net (represented by a grid) which functions as a transducer from physical signals to 2-D topographic maps, and an SRN which learns sequences of topographic maps corresponding to time-sliced acoustic and visual inputs respectively. Between them is the link SRN (2), which takes the acoustic SRN hidden layer configuration as input and the corresponding visual SRN hidden layer configuration as target output. The link SRN hidden layer configuration sequence represents the association of acoustic and visual input sequences.

This addresses (i) and (ii) above as follows:

i. Given an appropriate physical environment, the net learns to implement the expression-to-meaning function on-line and incrementally. There is no need to predefine an (expression, meaning) list.

ii. The net generates its own representations: expression representations are abstracted from acoustic input in the hidden layer of the acoustic SRN, and meaning representations develop in the link SRN hidden layer on the course of training; the intention is that, at any stage of acoustic input, the link net hidden layer should represent the meaning of the string up to that point, where 'meaning' is understood in an impoverished sense appropriate to the restricted range of inputs.

Conclusion

This paper has argued that pessimistic assessments of the adequacy of recurrent ANNs for NLP on the grounds that they have a finite state architecture are unjustified, and that their adequacy in this regard is an empirical issue. Part 1 presented arguments countering standard objections to finite state NLP on the grounds that these objections confuse the explanatory aims of linguistic theory with the essentially technological aims of NLP. Part 2 then went on to propose a finite state NLP model which maps strings onto meaning representations, presented preliminary string processing test and cluster analysis results of a computer simulation, and briefly addressed some problems and further developments. Whether or not the model can be developed for general NLP remains to be seen, but even as it stands it exemplifies the kind of radical departure from linguistics-based NLP that is possible once the supposed theoretical obstacles to finite state NLP are removed. In particular, it departs from linguistics-based NLP in making no use of any syntactic or compositional structure beyond the purely sequential, and amounts to table lookup mapping from strings to meaning representations.

References

Arbib, M. (1987) *Brains, machines, and mathematics*, 2nd ed. Springer, New York.

Blank, D., Meeden, L. & Marshall, J. (1992) 'Exploring the symbolic/subsymbolic continuum: a case study of RAAM', in: *The symbolic and connectionist paradigms: closing the gap*, ed. J. Dinsmore. Lawrence Erlbaum.

Cann, R. (1993) *Formal semantics: an introduction*. Cambridge University Press, Cambridge UK.

Castano, M., Vidal, E. & Casacumberta, F. (1993) 'Inference of stochastic regular languages through simple recurrent networks', in Grammatical Inference: theory, applications, and alternatives, ed. S. Lucas, IEE, London.

Chalmers, D. (1990) 'Syntactic transformations on distributed representations', *Connection Science* 2, 53-62.

Chomsky, N. (1956) 'Three models for the description of language', *IRE Transactions on Information Theory* IT-2, 109.

Church, K. (1980) 'On memory limitations in natural language processing', Tech. Report MIT/LCS/TR-45.

Dinsmore, J. (1992) *The symbolic and connectionist paradigms: closing the gap* Lawrence Erlbaum.

Elman, J. (1990) 'Finding structure in time', *Cognitive Science* 14, 179-211.

Elman, J. (1991) 'Distributed representation, simple recurrent networks, and grammatical structure', *Machine Learning* 7, 195-225.

Fodor, J. & Pylyshyn, Z. (1988) 'Connectionism and cognitive science: a critical analysis', *Cognition* 28, 3-71.

Gazdar, G. & Mellish, C. (1989) *Natural language processing in LISP*. Addison-Wesley.

Giles, C., Sun, G., Chen, Lee, Y, & Chen, D. (1990) 'Higher order recurrent networks and grammatical inference', in: *Advances in neural infromation systems 2*, ed. D. Touretzky. Morgan Kaufmann.

Giles, C., Chen, D., Miller, C., Chen, H., Sun, G., & Lee, Y. (1991) Grammatical inference using second-order recurrent neural networks, *Proceedings of the International Joint Conference on Neural Networks*, IEEE 91 CH3049-4.

Giles, C., Miller, C., Chen, D., Sun, G., Chen, H. & Lee, Y. (1992) 'Extracting and learning an unknown grammar with recurrent neural networks', in: *Advances in neural information processing systems*, ed. J. Moody, S. Hanson, R. Lippmann. Morgan Kaufmann.

Giles, C. & Omlin, C. (1993) 'Extraction, insertion, and refinement of symbolic rules in dynamically driven recurrent neural networks', *Connection Science* 5, 307-338.

Hopcroft, J. & Ullman, J. (1979) *Introduction to automata theory, languages, and computation*. Addison-Wesley.

Miller, G. & Chomsky, N. (1965) 'Finitary models of language users', in: *Readings in mathematical psychology II*, ed. R. Bush, E. Galanter, D. Luce. John Wiley.

Peschl, M. (1992) 'Construction, Representation, and the Embodiment of Knowledge, Meaning, and Symbols in Neural Structures, *Connection Science* 4, 327-338.

Pfeifer, R. & Verschure, P. (1992) 'Beyond Rationalism: Symbols, Patterns, and Behaviour, *Connection Science* 4, 313-325.

Plunkett, K., Sinha, C., Moeller, M. & Strandsby, O. (1992) 'Symbol Grounding or the Emergence of Symbols?', *Connection Science* 4, 293-312.

Pollack, J. (1990) 'Recursive distributed representations', *Artificial Intelligence* 46, 77-105.

Pollack, J. (1991) 'The induction of dynamical recognizers', *Machine Learning* 7, 123-148.

Pulman, S. (1986) 'Grammars, parsers, and memory limitations', *Language and Cognitive Processes* 1, 197-225.

Sanfeliu, A. & Alquezar, R. (1992) 'Understanding neural networks for grammatical inference and recognition', in: *Advances in structural and syntactic pattern recognition*, ed. H. Bunke. World Scientific.

Schwarz, G. (1992) 'Connectionism, processing, memory', *Connection Science* 4, 207-226.

Servan-Schreiber, D., Cleeremans, A. & McClelland, J. (1989) 'Learning sequential structure in simple recurrent networks', in: *Advances in neural information processing systems 1*, ed. D. Touretzky. Morgan Kaufmann.

Servan-Schreiber, D., Cleeremans, A. & McClelland, J. (1991) 'Graded state machines: the representation of temporal contingencies in simple recurrent networks', *Machine Learning* 7, 161-193

Sharkey, N. & Reilly, R. (1992) *Connectionist approaches to natural language processing*. Lawerence Erlbaum.

Sharkey, A. & Sharkey, N. (1993) 'Connectionism and natural language', in: *Grammatical inference: theory, applications, and alternatives*, ed. S. Lucas. IEE, London.

Smolensky, P. (1990) 'Tensor product variable binding and the representation of symbolic structures in connectionist systems', *Artificial Intelligence* 46, 159-216.

van Gelder, T. (1990) 'Compositionality: a connectionist variation on a classical theme', *Cognitive Science 14*, 355-384.

Watrous, R. & Kuhn, G. (1992) 'Induction of finite state languages using second-order recurrent networks', *Neural Computation* 4, 406-414.